Adjustment with persistent noise

Drew Fudenberg & Sergei Izmalkov

November 3, 2003

1 Motivation

- Select among strict Nash equilibria, some will not be "stochastically stable." E.g., in 2×2 games the risk-dominant equilibrium is selected (not necessarily Pareto efficient).
- Models are stochastic even in the limit (shocks do not diminish in time).
- There will always be small but positive probability of mutation (unlike ESS).
- A more appropriate model if mutations are real and recurring.
- Numerous applications. E.g., it is possible to incorporate mistakes in the learning process.

2 Stochastic Adjustment (KMR, Young, 93)

General procedure:

1. *Specify the game.* (will follow KMR)

- 2×2 symmetric game. (A, B) – actions.
- Suppose there are 3 NE:

  (A, A); (B, B); $(\alpha^* A + (1-\alpha^*)B, \alpha^* A + (1-\alpha^*)B)$.
- Suppose $\alpha^* < \frac{1}{2} \Rightarrow (A, A)$ – risk-dominant NE.

	A	B
A	2,2	0,0
B	0,0	1,1

Here $\alpha^* = \frac{1}{3}$.

2. *Specify a state space Θ (finite), e.g., the $\#s$ of players playing each strategy.*

- $\theta_t \in \Theta = [0, \ldots, N]$ – $\#$ of players using A.
- Denote

  $$u_A(\theta_t) = \frac{\theta_t}{N} u(A, A) + \frac{N - \theta_t}{N} u(A, B);\ u_B(\theta_t) = \ldots.$$

Possible to consider individual states (different states of the same agent), past histories,....

Steps 1 and 2 describe the game and provide structure for the analysis.

3. *Specify an "intentional" or "unperturbed" adjustment dynamics*, with a transition matrix P, where

$$P_{\theta,\xi} = \text{Pr}(\theta \text{ at } t+1 | \xi \text{ at } t).$$

- "Darwinian" dynamics: $\theta_{t+1} = P(\theta_t)$, where

$$sgn(P(\theta_t) - \theta_t) = sgn(u_A(\theta_t) - u_B(\theta_t)).$$

  Extreme: Best-response dynamics:

$$\theta_{t+1} = BR(\theta_t) = \begin{cases} N, & \text{for } u_A(\theta_t) > u_B(\theta_t), \\ \theta_t, & \text{for } u_A(\theta_t) = u_B(\theta_t), \\ 0, & \text{for } u_A(\theta_t) < u_B(\theta_t). \end{cases}$$

- Do we have learning here?

A simpler (different) example: (simultaneous move Cournot adjustment, Canning '92).

- If $N = 2$ (one player of each type), then can write:

$$P = \begin{array}{c c} \begin{array}{cccc} \text{AA} & \text{AB} & \text{BA} & \text{BB} \end{array} & \\ \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} & \begin{array}{c} \text{AA} \\ \text{AB} \\ \text{BA} \\ \text{BB} \end{array} \end{array}$$

Note a two-cycle between (AB) and (BA).

4. *Introduce a small noise: mistakes, mutations, etc...*

Consider P^{ε}, continuous in ε and $P^{\varepsilon} \to P$ as $\varepsilon \to 0$.

Make sure that P^{ε} is ergodic:

(a) There exist a unique (invariant distribution) ϕ^* (a probability distribution, a column vector) s.t.

$$\phi^* = \Pi\phi^*.$$

(b) Time averages converge

$$\text{a.s.} \lim_{T\to\infty} \left(\frac{1}{T}\right) \sum_t I(\theta_t = \theta) = \phi^*(\theta).$$

(c) Date-t distributions converge

$$\forall\phi, \lim_{t\to\infty} \Pi\phi = \phi^*.$$

Examples: $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

After perturbation?

Sufficient conditions for ergodicity of Π:

- $\Pi > 0$, or $\Pi^n > 0$ for some n.
- $\exists\theta$, that $\Pi_{\theta\theta} > 0$, and is reached from any other state for some n. (think recurrent classes)

In KMR, step 4.

- 2ε – probability that a player "mutates" (is replaced) (*after her intended choice), independent across players.

- Note: even if only 1 player "consciously" adjusts at a time, there is a positive probability that the whole population mutates at once.

- Clearly P^{ε} is ergodic.

5. *Verify that* $\lim_{\varepsilon \to 0} \phi^*_{\varepsilon} = \phi^*$ *exists*; compute ϕ^*. (By continuity $\phi^* = P\phi^*$.)

6. *Check that* ϕ^* *is a point mass*, i.e.,

$$\phi^*(\theta^*) = 1$$

for some θ^*.

The strategy profile at θ^* is called *stochastically stable equilibrium* (Foster & Young '90).

In the (simple) example :

$$P^{\varepsilon} = \begin{array}{cc} & \begin{array}{cccc} \text{AA} & \text{AB} & \text{BA} & \text{BB} \end{array} \\ & \left(\begin{array}{cccc} (1-\varepsilon)^2 & (1-\varepsilon)\varepsilon & (1-\varepsilon)\varepsilon & \varepsilon^2 \\ (1-\varepsilon)\varepsilon & \varepsilon^2 & (1-\varepsilon)^2 & (1-\varepsilon)\varepsilon \\ (1-\varepsilon)\varepsilon & (1-\varepsilon)^2 & \varepsilon^2 & (1-\varepsilon)\varepsilon \\ \varepsilon^2 & (1-\varepsilon)\varepsilon & (1-\varepsilon)\varepsilon & (1-\varepsilon)^2 \end{array} \right) \begin{array}{c} \text{AA} \\ \text{AB} \\ \text{BA} \\ \text{BB} \end{array} \end{array}$$

$$\phi^*_{\varepsilon} = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)^T.$$

Not an equilibrium, not a description of play in each period.

(Stable) cycles is the problem...

2.1 Limiting distribution (in Ex^0:BR)

- N^* is $\arg\min_m (m > N\alpha^*)$;
- $BR(\theta_t \geq N^*) = A$;
- $D_A = \{\theta \geq N^*\}$, $D_B = \{\theta < N^*\}$.
- Only basins of attraction matter: Intentional play depends on which of the two states, D_A or D_B, θ_t is and not on θ_t itself.

3 Result

Proposition: If N is large enough so that $N^* < \frac{N}{2}$, then limit φ^* of invariant distributions puts a point mass on $\theta_t = N$, corresponding to all players playing A.

Proof:

1. For any $\theta_t \in D_A$ ($\in D_B$) probability distribution $P^\varepsilon(\theta_t)$ is the same (think about intended choices and mistakes separately) — the problem can be reduced to two states.

2. Define

$$\begin{aligned} q_{BA} &= \Pr(\theta_{t+1} \in D_B | \theta_t \in D_A); \\ q_{AB} &= \Pr(\theta_{t+1} \in D_A | \theta_t \in D_B). \end{aligned}$$

3. Solve

$$\begin{bmatrix} \varphi_1 \\ \varphi_2 \end{bmatrix} = \begin{bmatrix} 1 - q_{AB} & q_{AB} \\ q_{BA} & 1 - q_{BA} \end{bmatrix} \begin{bmatrix} \varphi_1 \\ \varphi_2 \end{bmatrix}$$

   and find

$$\frac{\varphi_2}{\varphi_1} = \frac{q_{BA}}{q_{AB}}.$$

4. Take $\lim_{\varepsilon \to 0}$ of $\frac{\varphi_2}{\varphi_1}$.

To change $A \to B$, at least $N - N^*$ mutations into B are needed; for $B \to A$ at least N^* mutations must happen:

$$\begin{aligned} q_{BA} &\approx \binom{N}{N^*} \varepsilon^{N-N^*} (1-\varepsilon)^{N^*}; \\ q_{AB} &\approx \binom{N}{N^*} \varepsilon^{N^*} (1-\varepsilon)^{N-N^*}. \end{aligned}$$

Thus $\frac{\varphi_2}{\varphi_1} \to 0$ as $\varepsilon \to 0$.

4 Summary

- Selection of risk-dominant equilibrium as the unique long-run steady-state in 2×2 games (almost all models).

- "Learning" procedures tend to select equilibria that are relatively robust to mutations — different from Pareto efficiency.

	A	B
A	2,2	$-a$,0
B	0,$-a$	1,1

  (B, B) is risk-dominant if $1 + a > 2$.

- Probabilities (ratios of them) of escaping basins of attraction matter.

5 Other Dynamics

ω-limit sets of deterministic process matter (those that will be visited infinitely often at least from some initial state).

- Freidlin and Wentzell '84 insight:

  Suffices to consider Markov process on ω-limit sets.

  In 2×2 population games with two strict equilibria, there are only two such sets, 0 and N, so the general case reduces to BR considerations.

- In general: count the minimal costs (number of mutations, e.g.) that are necessary to pay to move from one ω-limit set to the basin of attraction of another.

6 Local interaction (Ellison)

- If the system starts near "wrong" equilibrium the expected time of adjustment may be quite large.

  In KMR model: the probability of escaping is $\approx \varepsilon^{N^*}$.

- Goal: to explain why stochastic adjustment processes might select the risk-dominant equilibrium in an economically relevant time frame.

- Players located on the circle and interact only with neighbors. A player selects an action and is matched randomly with one of the two neighbors (further extended).

- Action selected is a best response to the observed play in the last period.

- Observation: Pair of adjacent As wins the population.

6.1 Adjustment process

1. 2×2 symmetric game. (A, B) – actions.

2. $\Theta = \{A, B\}^N$.

3. Deterministic process: player with A switches its neighbors to A.

   Steady states: "All A", "All B", "$ABAB\ldots - BABA\ldots$" cycle.

4. Noise: Probability 2ε of mutating.

5. Limiting distribution: "All A",

   Convergence: Minimal cost of transition from "all B" is 2 if N is even and is 1 if N is odd. (number of mutations it takes to switch to "all A".)